arXiv:1608.07739v2 [cs.LG] 27 Feb 2017

Bayesian selection for the ℓ_2-Potts model regularization parameter: 1D piecewise constant signal denoising

Jordan Frecon*, Nelly Pustelnik*,
Nicolas Dobigeon†, Herwig Wendt†, and Patrice Abry*

Abstract

Piecewise constant denoising can be solved either by deterministic optimization approaches, based on the Potts model, or by stochastic Bayesian procedures. The former lead to low computational time but require the selection of a regularization parameter, whose value significantly impacts the achieved solution, and whose automated selection remains an involved and challenging problem. Conversely, fully Bayesian formalisms encapsulate the regularization parameter selection into hierarchical models, at the price of high computational costs. This contribution proposes an operational strategy that combines hierarchical Bayesian and Potts model formulations, with the double aim of automatically tuning the regularization parameter and of maintaining computational efficiency. The proposed procedure relies on formally connecting a Bayesian framework to a ℓ_2-Potts functional. Behaviors and performance for the proposed piecewise constant denoising and regularization parameter tuning techniques are studied qualitatively and assessed quantitatively, and shown to compare favorably against those of a fully Bayesian hierarchical procedure, both in accuracy and in computational load.

*J. Frecon, N. Pustelnik (Corresponding author) and P. Abry are with Univ Lyon, Ens de Lyon, Univ Claude Bernard, CNRS, Laboratoire de Physique, F-69342 Lyon, France (email: firstname.lastname@ens-lyon.fr).

†N. Dobigeon and H. Wendt are with IRIT/INP-ENSEEIHT, University of Toulouse, France (email: firstname.lastname@irit.fr).

1 Introduction

Piecewise constant denoising. Piecewise constant denoising (tightly related to change-point detection) is of considerable potential interest in numerous signal processing applications including, e.g., econometrics and biomedical analysis (see [1, 2], for an overview and [3] for an interesting application in biology). An archetypal and most encountered in the literature formulation considers noisy observations as resulting from the additive mixture of a piecewise constant signal $\overline{\boldsymbol{x}} \in \mathbb{R}^N$ with a Gaussian noise $\boldsymbol{\epsilon} \in \mathcal{N}(0, \sigma^2 \boldsymbol{I}_N)$

$$\boldsymbol{y} = \overline{\boldsymbol{x}} + \boldsymbol{\epsilon}. \tag{1}$$

Detecting change-points or denoising the piecewise constant information has been addressed by several strategies, such as Cusum procedures [1], hierarchical Bayesian inference frameworks [4, 5], or functional optimization formulations, involving either the ℓ_1-norm [6–10] or the ℓ_0-pseudo-norm of the first differences of the signal [11–14].

This latter class frames the present contribution. Formally, it amounts to recovering the solution of a ℓ_2-Potts model, namely,

$$\widehat{\boldsymbol{x}}_\lambda = \arg\min_{\boldsymbol{x} \in \mathbb{R}^N} \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \lambda \|\boldsymbol{L}\boldsymbol{x}\|_0, \tag{2}$$

where $\boldsymbol{L} \in \mathbb{R}^{(N-1)\times N}$ models the first difference operator, i.e., $\boldsymbol{L}\boldsymbol{x} = (x_{i+1} - x_i)_{1 \leq i \leq N-1}$, the ℓ_0-pseudo-norm counts the non-zero elements in $\boldsymbol{L}\boldsymbol{x}$, and $\lambda > 0$ denotes the regularization parameter which adjusts the respective contributions of the data-fitting and penalization terms in the objective function.

Such a formulation however suffers from a major limitation: its actual solution depends drastically on the regularization parameter λ. The challenging question of automatically estimating λ from data constitutes the core issue addressed in the present contribution.

Related works. Bayesian hierarchical inference frameworks have received considerable interests for addressing change-point or piecewise denoising problems [4,5]. This mostly results from their ability to include the hyperparameters within the Bayesian modeling and to jointly estimate them with the parameters of interest. In return for this intrinsic flexibility, approximating the Bayesian estimators associated with this hierarchical model generally requires the use of Markov chain Monte Carlo (MCMC) algorithms, which are often known as excessively demanding in terms of computational burden.

Remaining in the class of deterministic functional minimization, the non-convexity of the objective function underlying (2) has sometimes been alleviated by a convex relaxation, i.e., the ℓ_0-pseudo-norm is replaced by the ℓ_1-norm

$$\widetilde{\boldsymbol{x}}_\tau = \arg\min_{\boldsymbol{x} \in \mathbb{R}^N} \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \tau \|\boldsymbol{L}\boldsymbol{x}\|_1. \tag{3}$$

In essence, such an approach preserves the same intuition of piecewise constant denoising, at the price of a shrinkage effect, but with the noticeable advantage of ensuring convexity of the resulting function to be minimized (see, e.g., [8] for an intermediate solution where the penalization term is not convex but where the global criterion stays convex). This ensures the convergence of the minimization algorithms [15–18] or straightforward computations [19, 20]. The formulation proposed in (3) has received considerable interest, because, besides the existence and performance of sound algorithmic resolution procedures, it can offer some convenient ways to handle the automated tuning of the regularization parameter $\tau > 0$. For instance, the *Stein unbiased risk estimate* (SURE) [21, 22] aims at producing an unbiased estimator that minimizes the mean squared error between $\overline{\boldsymbol{x}}$ and $\widetilde{\boldsymbol{x}}_\tau$. While practically effective, implementing SURE requires the prior knowledge of the variance σ^2 of the residual error $\boldsymbol{\epsilon}$, often unavailable a priori (see, a contrario, [23, 24] for hyperspectral denoising or image deconvolution involving frames where σ^2 has been estimated). In [25], the regularization parameter τ is selected according to an heuristic rule, namely $\tau = 0.25\sqrt{N}\sigma$, derived in [8].

Alternatively, again, the problem in (3) can be tackled within a fully Bayesian framework, relying on the formulation of (3) as a statistical inference problem. Indeed, in the right-hand side of (3), the first term can be straightforwardly associated with a negative log-likelihood function by assuming an additive white Gaussian noise sequence $\boldsymbol{\epsilon}$, i.e., $\boldsymbol{y}|\boldsymbol{x}$ is distributed according to the Gaussian distribution $\mathcal{N}(\boldsymbol{x}, \sigma^2 \boldsymbol{I}_N)$. Further, the second term refers to a Laplace prior distribution for the first difference $\boldsymbol{L}\boldsymbol{x}$ of the unobserved

signal. Under such Bayesian modeling, the corresponding maximum a posteriori (MAP) criterion reads

$$\underset{\boldsymbol{x}\in\mathbb{R}^N}{\text{maximize}} \left\{ \left(\frac{1}{2\pi\sigma^2} \right)^{N/2} e^{-\frac{1}{2\sigma^2}\|\boldsymbol{y}-\boldsymbol{x}\|_2^2} \frac{1}{Z(\tau/\sigma^2)} e^{-\frac{\tau}{\sigma^2}\|\boldsymbol{L}\boldsymbol{x}\|_1} \right\} \tag{4}$$

whose resolution leads to the solution (3) and where $Z(\tau/\sigma^2)$ is the normalizing constant associated with the prior distribution. Following a hierarchical strategy, the hyperparameters τ and σ^2 could be included into the Bayesian model to be jointly estimated with $\boldsymbol{x}$. However, in the specific case of (3), the prior distribution related to the penalization is not separable with respect to (w.r.t.) the individual components of $\boldsymbol{x}$: The partition function $Z(\tau/\sigma^2)$ can hence not be expressed analytically. Therefore, estimating τ within a hierarchical Bayesian framework would require either to choose a heuristic prior for τ as proposed in [26–28] or to conduct intensive approximate Bayesian computations as in [29].

Goals, contributions and outline. Departing from the Bayesian interpretation of (3), the present contribution chooses to stick to the original ℓ_2-Potts model (2). Capitalizing on efficient dynamic programming algorithms [14, 30–33] which allow $\widehat{\boldsymbol{x}}_\lambda$ to be recovered for a predefined value of λ, the main objective of this work resides in the joint estimation of the denoised signal and the *optimal* hyperparameter λ, without assuming any additional prior knowledge regarding the residual variance σ^2. Formally, this problem can be formulated as an extended counterpart of (2), i.e., a minimization procedure involving $\boldsymbol{x}$, λ, and σ^2 as stated in what follows.

Problem 1.1 *Let $\boldsymbol{y} \in \mathbb{R}^N$ and $\phi\colon \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$. The problem under consideration is*[1]

$$\underset{\boldsymbol{x}\in\mathbb{R}^N,\lambda>0,\sigma^2>0}{\text{minimize}} \frac{1}{2\sigma^2}\|\boldsymbol{y}-\boldsymbol{x}\|_2^2 + \frac{\lambda}{\sigma^2}\|\boldsymbol{L}\boldsymbol{x}\|_0 + \phi(\lambda,\sigma^2). \tag{5}$$

The main challenge for handling Problem 1.1 lies in the design of an appropriate function ϕ that leads to a relevant penalization of the overall criterion w.r.t. the set of nuisance parameters (λ, σ^2). To that end, Section 2 provides a natural parametrization of $\boldsymbol{x}$ and a reformulation of Problem 1.1. In Section 3, a closed-form expression of $\phi(\cdot,\cdot)$ (cf. (28)) and an interpretation of λ will be derived from a relevant hierarchical Bayesian inference framework. For this particular choice of the function $\phi(\cdot,\cdot)$, Section 4 proposes an efficient algorithmic strategy to approximate a solution of the Problem 1.1. In Section 5, the relevance and performance of this procedure are qualitatively illustrated and quantitatively assessed, and shown to compare favorably against the Markov chain Monte Carlo algorithm resulting from the hierarchical Bayesian counterpart of (5), both in terms of accuracy and in computational load.

2 Problem parametrization

Following [4, 5, 34–36], piecewise constant signals $\boldsymbol{x} \in \mathbb{R}^N$ can be explicitly parametrized via change-point locations $\boldsymbol{r}$ and amplitudes of piecewise constant segments $\boldsymbol{\mu}$. These reparametrizations are derived in Sections 2.1 and 2.2. They are in turn used in Section 2.3 to bring Problem 1.1 in a form more amenable for explicit connection to a hierarchical Bayesian model.

2.1 Change-point location parametrization $\boldsymbol{r}$

To locate the time instants of the change-points in the denoised signal $\boldsymbol{x}$, an indicator vector $\boldsymbol{r} = \big(r_i\big)_{1\leq i\leq N} \in \{0,1\}^N$ is introduced as follows

$$r_i = \begin{cases} 1, & \text{if there is a change-point at time instant } i, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

[1]Note that (5) could have been normalized differently without changing the minimization problem. A usual formulation aims at multiplying all terms by σ^2. For our study, formulation (5) is adopted for convenience.

By convention, $r_i = 1$ indicates that x_i is the last sample belonging to the current segment, and thus that x_{i+1} belongs to the next segment. Moreover, stating $r_N = 1$ ensures that the number K of segments is equal to the number of change-points, i.e., $K = \sum_{i=1}^{N} r_i$.

For each segment index $k \in \{1, \dots, K\}$, the set $\mathcal{R}_k \subset \{1, \dots, N\}$ is used to denote the set of time indices associated with the k-th segment. In particular, it is worthy to note that $\mathcal{R}_k \cap \mathcal{R}_{k'} = \emptyset$ for $k \neq k'$ and $\cup_{k=1}^{K} \mathcal{R}_k = \{1, \cdots, N\}$. Hereafter, the notation $K_{\boldsymbol{r}}$ will be adopted to emphasize the dependence of the number K of segments on the indicator vector $\boldsymbol{r}$, i.e., $K = \|\boldsymbol{r}\|_0$.

2.2 Segment amplitude parametrization $\boldsymbol{\mu}$

The amplitudes of each segment of the piecewise constant signal can be encoded by introducing the vector $\boldsymbol{\mu} = (\mu_k)_{1 \leq k \leq K_{\boldsymbol{r}}}$ such that

$$(\forall k \in \{1, \dots, K_{\boldsymbol{r}}\})(\forall i \in \mathcal{R}_k) \quad x_i = \mu_k. \tag{7}$$

2.3 Reformulation of Problem 1.1

In place of $\boldsymbol{x}$, the parameter vector $\boldsymbol{\theta} = \{\boldsymbol{r}, \boldsymbol{\mu}\}$ will now be used to fully specify the piecewise constant signal $\boldsymbol{x}$. An important issue intrinsic to the ℓ_2-Potts model and thus to this formulation stems from the fact that the unknown parameter $\boldsymbol{\theta}$ belongs to the space

$$\mathcal{S} = \{\{0,1\}^N \times \mathbb{R}^{K_{\boldsymbol{r}}} : K_r = \{1, ..., N-1\}\}$$

whose dimension is a priori unknown, as it depends on the number $K_{\boldsymbol{r}}$ of change-points. Moreover, this parametrization leads to the following lemma.

Lemma 2.1 *Let $\boldsymbol{y} \in \mathbb{R}^N$ and $\phi \colon \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$. Problem 1.1 is equivalent to*

$$\underset{\substack{\boldsymbol{\theta}=\{\boldsymbol{r},\boldsymbol{\mu}\}\in\mathcal{S} \\ \lambda>0,\sigma^2>0}}{\text{minimize}} \; \frac{1}{2\sigma^2} \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - \mu_k)^2 + \frac{\lambda}{\sigma^2}(K_{\boldsymbol{r}} - 1) + \phi(\lambda, \sigma^2) \tag{8}$$

where $(\mathcal{R}_k)_{1 \leq k \leq K_{\boldsymbol{r}}}$ is related to $\boldsymbol{r}$ as indicated in Section 2.1.

Indeed, the data fidelity term in the minimization Problem 1.1 can be equivalently written as

$$\|\boldsymbol{y} - \boldsymbol{x}\|^2 = \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - x_i)^2 = \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - \mu_k)^2. \tag{9}$$

Moreover, the penalization can be rewritten as

$$\|\boldsymbol{L}\boldsymbol{x}\|_0 = \|\boldsymbol{r}\|_0 - 1 = K_{\boldsymbol{r}} - 1. \tag{10}$$

Lemma 2.1 implies that estimating the piecewise constant signal $\boldsymbol{x}$ can be equivalently formulated as estimating the parameter vector $\boldsymbol{\theta}$.

3 Bayesian derivation of ϕ

Assisted by the reformulation of Problem 1.1 and a hierarchical Bayesian framework detailed in Section 3.1, a relevant penalization function ϕ will be derived in Section 3.2.

3.1 Hierarchical Bayesian model

In [4, 5], the problem of detecting change-points in a stationary sequence has been addressed following a Bayesian inference procedure which aims at deriving the posterior distribution of the parameter vector $\boldsymbol{\theta}$ from the likelihood function associated with the observation model and the prior distributions chosen for the unknown parameters. In what follows, a similar approach is proposed to produce a hierarchical Bayesian model that can be tightly related to the Problem 1.1 under a joint MAP paradigm.

First, the noise samples $(\boldsymbol{\epsilon}_i)_{1\leq i\leq N}$ are assumed to be independent and identically distributed (i.i.d.) zero mean Gaussian variables with common but unknown variance σ^2, i.e., $\boldsymbol{\epsilon}|\sigma^2 \sim \mathcal{N}\left(\mathbf{0}, \sigma^2 \boldsymbol{I}_N\right)$. The resulting joint likelihood function of the observations $\boldsymbol{y}$ given the piecewise constant model $\{\boldsymbol{r}, \boldsymbol{\mu}\}$ and the noise variance σ^2 reads

$$f\left(\boldsymbol{y}|\boldsymbol{r}, \boldsymbol{\mu}, \sigma^2\right) = \prod_{k=1}^{K_{\boldsymbol{r}}} \prod_{i\in\mathcal{R}_k} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\left(\mu_k - y_i\right)^2}{2\sigma^2}\right). \tag{11}$$

Then to derive the posterior distribution, prior distributions are elected for the parameters $\boldsymbol{r}$ and $\boldsymbol{\mu}$, assumed to be a priori independent. Following well-admitted choices such as those in [4, 5, 34, 35, 37], the components r_i of the indicator vector $\boldsymbol{r}$ are assumed to be a priori independent and identically distributed (i.i.d.) according to a Bernoulli distribution of hyperparameter p

$$f(\boldsymbol{r}|p) = \prod_{i=1}^{N-1} p^{r_i}(1-p)^{1-r_i} = p^{\sum_{i=1}^{N-1} r_i}(1-p)^{(N-1-\sum_{i=1}^{N-1} r_i)} = \left(\frac{p}{1-p}\right)^{(K_{\boldsymbol{r}}-1)}(1-p)^{(N-1)}. \tag{12}$$

The prior independence between the indicator components r_i $(i = 1, \ldots, N-1)$ implicitly assumes that the occurrence of a change at a given time index does not depend on the occurrence of change at any other time index. Moreover, the hyperparameter p stands for the prior probability of occurrence of a change, which is assumed to be independent of the location. Obviously, for particular applications, alternative and more specific choices can be adopted relying, e.g., on hard constraints [38] or Markovian models [39], for instance to prevent solutions composed of too short segments.

From a Bayesian perspective, a natural choice for $f(\boldsymbol{\mu}|\boldsymbol{r})$ consists in electing independent conjugate Gaussian prior distributions $\mathcal{N}\left(\mu_0, \sigma_0^2\right)$ for the segment amplitudes μ_k $(k = 1, \ldots, K_{\boldsymbol{r}})$, i.e.,

$$f(\boldsymbol{\mu}|\boldsymbol{r}) = \prod_{k=1}^{K_{\boldsymbol{r}}} \frac{1}{\sqrt{2\pi\sigma_0^2}} e^{\frac{-(\mu_k-\mu_0)^2}{2\sigma_0^2}}. \tag{13}$$

Indeed, this set of conjugate priors ensures that the conditional posterior distributions of the segment amplitudes are still Gaussian distributions.

Moreover, within a hierarchical Bayesian paradigm, nuisance parameters, such as the noise variance, and other hyperparameters defining the prior distributions can be included within the model to be estimated jointly with $\boldsymbol{\theta}$ [4, 5]. In particular, to account for the absence of prior knowledge on the noise variance σ^2, a non-informative Jeffreys prior is assigned to σ^2

$$f\left(\sigma^2\right) \propto \frac{1}{\sigma^2}. \tag{14}$$

Proposed in [40], the use of this improper distribution has been widely advocated in the Bayesian literature for its invariance under reparametrization [41, Chap. 3] (see also [35, 37]). Finally, as in [5, 35, 36, 39], a conjugate Beta distribution $\mathcal{B}(\alpha_0, \alpha_1)$ is assigned to the unknown hyperparameter p, which is a natural choice to model a $(0, 1)$-constrained parameter

$$f(p) = \frac{\Gamma\left(\alpha_0 + \alpha_1\right)}{\Gamma\left(\alpha_0\right)\Gamma\left(\alpha_1\right)} p^{\alpha_1 - 1}(1-p)^{\alpha_0 - 1}. \tag{15}$$

Note that a wide variety of distribution shapes can be obtained by tuning the two hyperparameters α_0 and α_1, while ensuring the parameter p belongs to the set $(0, 1)$, as required since p stands for a probability [41, App. A]. In particular, when the hyperparameters are selected as $\alpha_0 = \alpha_1 = 1$, the prior in (15) reduces to the uniform distribution.

3.2 Joint MAP criterion

From the likelihood function (11) and prior and hyper-prior distributions (12)–(15) introduced above, the joint posterior distribution reads

$$f(\boldsymbol{\Theta}|\boldsymbol{y}) \propto\ f\left(\boldsymbol{y}|\boldsymbol{r},\boldsymbol{\mu},\sigma^2\right) f\left(\boldsymbol{\mu}|\boldsymbol{r}\right) f(\boldsymbol{r}|p) f(p) f\left(\sigma^2\right) \tag{16}$$

with $\boldsymbol{\Theta} = \left\{\boldsymbol{r},\boldsymbol{\mu},\sigma^2,p\right\}$. Deriving the Bayesian estimators, such as the minimum mean square error (MMSE) and MAP estimators associated with this posterior distribution is not straightforward, mainly due to the intrinsic combinatorial problem resulting from the dimension-varying parameter space $\{0,1\}^N \times \mathbb{R}^{K_{\boldsymbol{r}}}$. In particular, a MAP approach would consist in maximizing the joint posterior distribution (16), which can be reformulated as the following minimization problem by taking the negative logarithm of (16).

Problem 3.1 *Let $\boldsymbol{y} = (y_i)_{1\leq i\leq N} \in \mathbb{R}^N$ and let $\boldsymbol{\Phi} = \left\{\alpha_0,\alpha_1,\sigma_0^2\right\}$ a set of hyperparameters. We aim to*

$$\begin{aligned}\underset{\boldsymbol{\Theta}=\{\boldsymbol{r},\boldsymbol{\mu},\sigma^2,p\}}{\text{minimize}}\ & \frac{1}{2\sigma^2}\sum_{k=1}^{K_{\boldsymbol{r}}}\sum_{i\in\mathcal{R}_k}(y_i-\mu_k)^2 + (K_{\boldsymbol{r}}-1)\left(\log\left(\frac{1-p}{p}\right)+\frac{1}{2}\log(2\pi\sigma_0^2)\right)\\ &+\frac{N}{2}\log(2\pi\sigma^2)-(N-1)\log(1-p)+\log\sigma^2-(\alpha_1-1)\log p-(\alpha_0-1)\log(1-p)\\ &+\frac{1}{2\sigma_0^2}\sum_{k=1}^{K_{\boldsymbol{r}}}(\mu_k-\mu_0)^2+\frac{1}{2}\log(2\pi\sigma_0^2).\end{aligned} \tag{17}$$

Despite apparent differences in parametrization between Problem 1.1 and Problem 3.1, we prove hereafter that both are equivalent for specific choices of λ and $\phi(\cdot,\cdot)$.

Proposition 3.1 *For σ_0^2 large enough, Problem 1.1 with*

$$\lambda = \sigma^2\left(\log\left(\frac{1-p}{p}\right)+\frac{1}{2}\log(2\pi\sigma_0^2)\right) \tag{18}$$

and

$$\begin{aligned}\phi(\lambda,\sigma^2) =& \frac{N}{2}\log(2\pi\sigma^2)+\log(\sigma^2)-\frac{\lambda}{\sigma^2}(N+\alpha_0-2)+\frac{N+\alpha_0-1}{2}\log(2\pi\sigma_0^2)\\ &+(N+\alpha_0+\alpha_1-3)\log\left(1+\exp\left(\frac{\lambda}{\sigma^2}-\frac{1}{2}\log(2\pi\sigma_0^2)\right)\right)\end{aligned} \tag{19}$$

matches Problem 3.1.

The sketch of the proof consists in identifying the three terms of the expression in (8) in the criterion (17): the data fidelity term (9), a term proportional to the regularization (10), and a third term $\phi(\lambda,\sigma^2)$ that is independent of the indicator vector $\boldsymbol{r}$. Identification is possible under the condition that the term $\frac{1}{2\sigma_0^2}\sum_{k=1}^{K_{\boldsymbol{r}}}(\mu_k-\mu_0)^2$ which explicitly depends on $\boldsymbol{r}$ through $K_{\boldsymbol{r}}$ can be neglected. Thus, choosing σ_0^2 sufficiently large

$$\frac{1}{2\sigma_0^2}\sum_{k=1}^{K_{\boldsymbol{r}}}(\mu_k-\mu_0)^2 \ll \frac{1}{2}\log(2\pi\sigma_0^2), \tag{20}$$

permits to equate Problem 1.1 and Problem 3.1 with the choices λ and $\phi(\lambda,\sigma^2)$ as defined in Proposition 3.1. As an illustration of its complex shape, Fig. 1 represents $\phi(\lambda,\sigma^2)$ as function of λ and σ^2.

Remark 1 *The tuning of σ_0^2 requires some discussion:*

- *As $\frac{1}{2\sigma_0^2}\sum_{k=1}^{K_{\boldsymbol{r}}}(\mu_k-\mu_0)^2$ is of the order of $\frac{1}{2\sigma_0^2}K_{\boldsymbol{r}}\sigma_0^2 \approx \frac{pN}{2}$, a sufficient condition for (20) to hold reads:*

$$\frac{pN}{2} \ll \frac{1}{2}\log(2\pi\sigma_0^2). \tag{21}$$

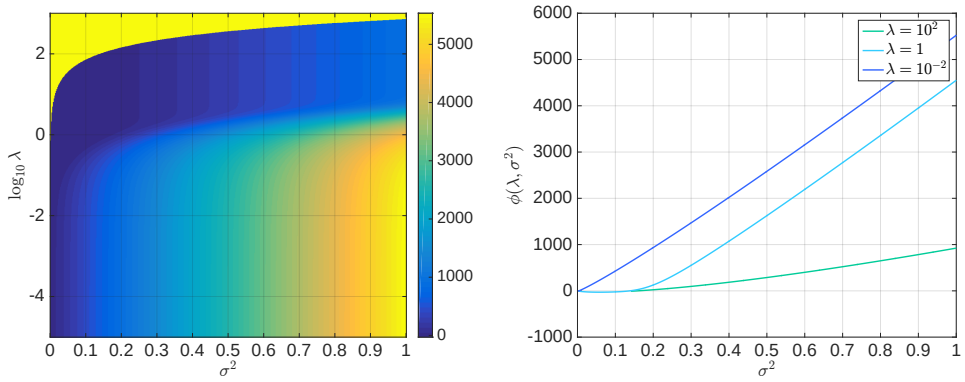

Figure 1: **Illustration of** $\phi(\lambda, \sigma^2)$ for the hyperparameter setting $\alpha_0 = \alpha_1 = 1$ and $2\pi\sigma_0^2 = 10^4$.

- *However a careful examination of* (17) *and* (18) *also leads to conclude that parameter* p *actually controls* λ *provided that* $\log\left(\frac{1-p}{p}\right)$ *is not totally neglectable when compared to* $\frac{1}{2}\log(2\pi\sigma_0^2)$*, thus implying an upper bound of the form:*

$$\frac{1}{2}\log(2\pi\sigma_0^2) \leq \log\frac{(1-p)}{p} + \log\zeta \tag{22}$$

where $\zeta > 0$ *is an arbitrary constant, whose magnitude will be precisely addressed in Section 5.5.*

The tuning of $\log(2\pi\sigma_0^2)$*, of paramount practical importance, is hence not intricate and will be further discussed Section 5.5 from numerical experimentations.*

4 Algorithmic solution

Thanks to Proposition 3.1, a function ϕ has been derived which allows the choice of the regularization parameter λ to be penalized in Problem 1.1. In this section, an algorithmic solution is proposed to estimate $(\widehat{\boldsymbol{x}}, \widehat{\lambda}, \widehat{\sigma}^2)$, a solution of Problem 1.1.

An alternate minimization over $\boldsymbol{x}$, λ and σ^2 would not be efficient due to the non-convexity of the underlying criterion. To partly alleviate this problem, we propose to estimate λ on a grid Λ. Therefore, a candidate solution can be obtained by solving $(\forall \lambda \in \Lambda)$

$$(\widehat{\boldsymbol{x}}_\lambda, \widehat{\sigma}_\lambda^2) \in \underset{\boldsymbol{x}\in\mathbb{R}^N, \sigma^2>0}{\operatorname{Argmin}} \; F(\boldsymbol{x}, \lambda, \sigma^2)$$

with

$$F(\boldsymbol{x}, \lambda, \sigma^2) = \frac{1}{2\sigma^2}\|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \frac{\lambda}{\sigma^2}\|L\boldsymbol{x}\|_0 + \phi(\lambda, \sigma^2). \tag{23}$$

The minimization over $\boldsymbol{x}$ does not depend on σ^2, thus

$$\widehat{\boldsymbol{x}}_\lambda = \arg\min_{\boldsymbol{x}\in\mathbb{R}^N} \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \lambda\|L\boldsymbol{x}\|_0, \tag{24}$$

and we set

$$\widehat{\sigma}_\lambda^2 = \frac{\|\boldsymbol{y} - \widehat{\boldsymbol{x}}_\lambda\|^2}{N-1}. \tag{25}$$

We finally select the triplet $(\widehat{\boldsymbol{x}}_{\widehat{\lambda}}, \widehat{\lambda}, \widehat{\sigma}_{\widehat{\lambda}}^2)$ such that

$$\widehat{\lambda} = \arg\min_{\lambda\in\Lambda} F(\widehat{\boldsymbol{x}}_\lambda, \lambda, \widehat{\sigma}_\lambda^2). \tag{26}$$

Note that the provided estimation amounts to using the solution of (2) for different $\lambda \in \Lambda$ to probe the space $(\boldsymbol{x}, \sigma^2) \in \mathbb{R}^N \times \mathbb{R}_+$. Therefore, the iterations of the proposed full algorithmic scheme (reported in

Algo. 1) are very succinct and the overall algorithm complexity mainly depends on the ability to solve (2) efficiently for any $\lambda \in \Lambda$. In this work, we propose to resort to a dynamic programming algorithm developed in [14, 31] that allows (24) to be solved exactly. We use its *Pottslab* implementation [32] augmented by a pruning strategy [42].

Algorithm 1 Bayesian driven resolution of the ℓ_2-Potts model

Input: Observed signal $\boldsymbol{y} \in \mathbb{R}^N$.
The predefined set of regularization parameters Λ.
Hyperparameters $\boldsymbol{\Phi} = \{\alpha_0, \alpha_1, \sigma_0^2\}$.

Iterations:

1: **for** $\lambda \in \Lambda$ **do**
2: Compute $\widehat{\boldsymbol{x}}_\lambda = \arg\min_{\boldsymbol{x}\in\mathbb{R}^N} \frac{1}{2}\|\boldsymbol{y}-\boldsymbol{x}\|_2^2 + \lambda\|\boldsymbol{L}\boldsymbol{x}\|_0$.
3: Compute $\widehat{\sigma}_\lambda^2 = \|\boldsymbol{y}-\widehat{\boldsymbol{x}}_\lambda\|^2/(N-1)$.
4: **end for**

Output: Solution $(\widehat{\boldsymbol{x}}_{\widehat{\lambda}}, \widehat{\lambda}, \widehat{\sigma}_{\widehat{\lambda}}^2)$ with $\widehat{\lambda} = \arg\min_{\lambda\in\Lambda} F(\widehat{x}_\lambda, \lambda, \widehat{\sigma}_\lambda^2)$

5 Automated selection of λ: Illustration and performance

5.1 Performance evaluation and hyperparameter settings

5.1.1 Synthetic data

The performance of the proposed automated selection of λ are illustrated and assessed using Monte Carlo numerical simulation based on synthetic data $\boldsymbol{y} = \overline{\boldsymbol{x}} + \boldsymbol{\epsilon}$, where the noise $\boldsymbol{\epsilon}$ consists of i.i.d. samples drawn from $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2\boldsymbol{I}_N)$, and the signal $\overline{\boldsymbol{x}}$ is piecewise constant, with i.i.d. change-points, occurring with probability p, and i.i.d. amplitudes drawn from a uniform distribution (on the interval $[\overline{x}_{\min}, \overline{x}_{\max}]$).

5.1.2 Performance quantification

Performance are quantified by the relative mean square error (MSE) and the Jaccard error. While the former evaluates performance in the overall (shape and amplitude) estimation $\widehat{\boldsymbol{x}}$ of $\overline{\boldsymbol{x}}$ such that

$$\mathrm{MSE}(\overline{\boldsymbol{x}}, \widehat{\boldsymbol{x}}) = \frac{\|\overline{\boldsymbol{x}} - \widehat{\boldsymbol{x}}\|}{\|\overline{\boldsymbol{x}}\|},$$

the latter focuses on the accuracy of change-point location estimation $\boldsymbol{r}$. The Jaccard error between the true change-point vector $\overline{\boldsymbol{r}}$ and its estimate $\widehat{\boldsymbol{r}}$ (both in $\{0,1\}^N$), is defined as [43, 44]

$$J(\overline{\boldsymbol{r}}, \widehat{\boldsymbol{r}}) = 1 - \frac{\sum_{i=1}^N \min(\overline{r}_i, \widehat{r}_i)}{\sum_{\substack{1\le i\le N \\ \overline{r}_i>0, \widehat{r}_i>0}} \frac{\overline{r}_i+\widehat{r}_i}{2} + \sum_{\substack{1\le i\le N \\ \widehat{r}_i=0}} \overline{r}_i + \sum_{\substack{1\le i\le N \\ \overline{r}_i=0}} \widehat{r}_i}. \tag{27}$$

$J(\overline{\boldsymbol{r}}, \widehat{\boldsymbol{r}})$ ranges from 0 when $\overline{\boldsymbol{r}} = \widehat{\boldsymbol{r}}$, to 1, when $\overline{\boldsymbol{r}} \cap \widehat{\boldsymbol{r}} = \emptyset$. The Jaccard error is a demanding measure: when one half of non-zero values of two given binary sequences coincides while the other half does not, then $J(\overline{\boldsymbol{r}}, \widehat{\boldsymbol{r}}) = 2/3$.

In the present study, to account for the fact that a solution with a change point position mismatch by a few time indices remains useful and of practical interest, the Jaccard error is computed between smoothed versions $\overline{\boldsymbol{r}} * \mathcal{G}$ and $\widehat{\boldsymbol{r}} * \mathcal{G}$ of the true and estimated sequences $\overline{\boldsymbol{r}}$ and $\widehat{\boldsymbol{r}}$. The convolution kernel $\mathcal{G}$ is chosen here as a Gaussian filter (with a stansdard deviation of 0.5) truncated to a 5-sample support.

Performance are averaged over 50 realizations, except for comparisons with the MCMC-approximated Bayesian estimators (see Section 5.4) where only 20 realizations are used because of MCMC procedure's high computational cost.

5.1.3 Hyperparameter setting

The prior probability p for change-point is chosen as a uniform distribution over $(0,1)$, obtained with hyper-parameters set to $\alpha_0 = \alpha_1 = 1$. Indeed, for such a choice, the Beta distribution in (15) reduces to a uniform distribution, hence leading to a non-informative prior for the change-point probability. This hyperparameter setting leads to the following expression for the penalization function

$$\phi(\lambda, \sigma^2) = \log(\sigma^2) + \frac{N}{2}\big(\log(2\pi\sigma^2) + \log(2\pi\sigma_0^2)\big) + (N-1)\left(\log\left(1 + \exp\left(\frac{\lambda}{\sigma^2} - \frac{1}{2}\log(2\pi\sigma_0^2)\right)\right) - \frac{\lambda}{\sigma^2}\right). \quad (28)$$

Amplitudes $\boldsymbol{\mu}$ for $\overline{\boldsymbol{x}}$ are parametrized with σ_0^2, which according to Proposition 3.1 should be chosen large enough. For the time being, we set $2\pi\sigma_0^2 = 10^4$, and further discuss the impact of this choice in Section 5.5.

5.1.4 Discretization of Λ

For practical purposes, we make use of a discretized subset Λ for λ (500 values equally spaced, in a $\log_{10}$-scale, between 10^{-5} and 10^5). Note that in the toolbox associated with this work, an option is called $\lambda-$shooting allowing to select the grid Λ according to the strategy introduced in [31].

5.2 Illustration of the principle of the automated tuning of λ.

Fig. 2 illustrates the principle of the automated selection of λ, under various scenarios, with different values for p and different *amplitude-to-noise-ratios*[2] (ANR) where ANR$= \frac{\overline{x}_{\max} - \overline{x}_{\min}}{3\sigma}$.

For all scenarios, Fig. 2 shows that the automatically selected $\widehat{\lambda}$, obtained as the minimum of the devised criterion F (cf. (23), vertical red line in bottom row), satisfactorily falls within the ranges of λ achieving the MSE minimum (denoted Λ_{MSE} and marked with vertical lines in the second row) or the minimum of Jaccard error (denoted Λ_{Jac} and marked with vertical lines in the third row): $\widehat{\lambda} \in \Lambda_{\mathrm{MSE}} \cap \Lambda_{\mathrm{Jac}}$. In addition, on the first row of Fig. 2, the corresponding solution $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ (red) visually appears as a satisfactory estimator of $\overline{\boldsymbol{x}}$ (black), similar to the "oracle" solutions $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{MSE}}}$ (blue) and $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{Jac}}}$ (green) that rely on a perfect knowledge of the noise-free signal $\overline{\boldsymbol{x}}$. Solution $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ indeed systematically benefits from lower relative MSE and Jaccard error than $\widehat{\boldsymbol{x}}_{\lambda}$ for any other λ. While, by construction, $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{MSE}}}$ and $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{Jac}}}$ are identical for all λ within Λ_{MSE} or Λ_{Jac}, the automated selection procedure for λ yields interestingly a single global minimizer.

When ANR decreases, a closer inspection of Fig. 2 (left column) further shows that the supports of oracle λ, Λ_{MSE} and Λ_{Jac} are drastically shrinking, yet the automated selection of λ remains satisfactory even in these more difficult contexts. The same holds when p increases (see Fig. 2, right column).

5.3 Estimation performance quantification

To assess and quantify estimation performance of $\widehat{\lambda}$ as functions of data parameters σ^2, $\overline{x}_{\max} - \overline{x}_{\min}$, and p, we have performed Monte Carlo simulations under various settings.

First, Fig. 3 reports estimation performance for $\widehat{\lambda}$ as a function of the ANR. It shows that the estimated value $\widehat{\lambda}$ (red), averaged over Monte Carlo simulations, satisfactorily remains within the range of MSE/Jaccard error oracle values for λ (dashed white lines) and tightly follows the average oracle values (solid white line). This holds for different $\overline{x}_{\max} - \overline{x}_{\min}$. As p grows (cf. Fig. 3 from top to bottom), the oracle regions in dashed white shrink, thus indicating that the selection of λ becomes more intricate when more segments are to be detected. The proposed automated selection for λ still performs well in these more difficult situations. In addition, it can also be observed that $\widehat{\lambda}$ depends, as expected, on σ (or equivalently on $\overline{x}_{\max} - \overline{x}_{\min}$) cf. Fig. 3, from left to right.

Second, Fig. 4 focuses on the behavior of the estimated $\widehat{\lambda}$ as a function of σ, for different values of ANR. Again, it shows satisfactory performance of $\widehat{\lambda}$ compared to the oracle λ. Incidently, it also very satisfactorily

[2]This measure allows amplitudes between successive segments to be compared w.r.t. to noise power. Since segment amplitudes are drawn uniformly between $\overline{x}_{\min}$ and $\overline{x}_{\max}$ the average difference between successive segments is $\frac{\overline{x}_{\max} - \overline{x}_{\min}}{3}$.

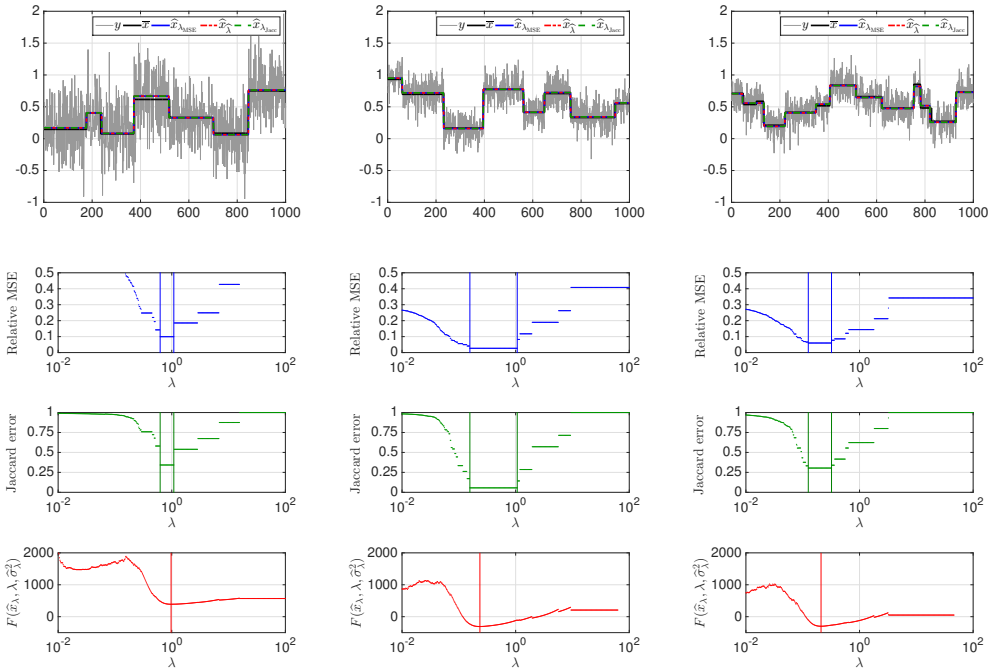

(a) ANR $= 1$, $p = 0.01$. (b) ANR $= 2$, $p = 0.01$. (c) ANR $= 2$, $p = 0.015$.

Figure 2: **Illustration of the automated tuning of λ:** Top: Data $\boldsymbol{y}$ to which are superimposed true signal $\overline{\boldsymbol{x}}$ and oracle signals $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{MSE}}}$ (blue) and $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{Jac}}}$ (green) obtained for λ_{MSE} and λ_{Jac} minimizing the MSE and the Jaccard error, together with estimated $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ (red) obtained from automated selection of λ. Second and third lines: relative MSE and Jaccard error as functions of λ. Vertical lines locate λ_{MSE} and λ_{Jac}. Bottom line: Criterion F (cf. (23)) as a function of λ. Automatically selected $\widehat{\lambda}$ is indicated by vertical red lines and is satisfactorily located in between the vertical lines indicating λ_{MSE} and λ_{Jac}. (left) $p = 0.01$ and ANR $= 1$, (middle) $p = 0.01$ and ANR $= 2$, (right) $p = 0.15$ and ANR $= 2$. For all configurations, $\overline{x}_{\max} - \overline{x}_{\min} = 1$.

reproduces the linear dependence of λ with σ^2, which can be predicted from a mere dimensional analysis of the ℓ_2-Potts model yielding:

$$\lambda \sim \frac{\sigma^2}{2p}. \tag{29}$$

5.4 Comparison with state-of-the-art estimators

The proposed method has been compared to classical Bayesian estimators associated with the hierarchical Bayesian model derived in Section 3.1 for which an MCMC procedure has been derived (cf. Appendix A). The number of Monte Carlo iterations is fixed to $T_{\mathrm{MC}} = 10^3$ and the amplitude hyperprior parameters are chosen as the mean of $\boldsymbol{y}$ for μ_0 and $\sigma_0^2 = \widehat{\mathrm{var}}(\boldsymbol{y})$, where $\widehat{\mathrm{var}}(\cdot)$ stands for the empirical variance.

In Fig. 5, estimation performance for the proposed procedure (solid red) are compared against MAP and MMSE hierarchical Bayesian estimators, as functions of ANR. Overall, $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ is equivalent to $\widehat{\boldsymbol{x}}_{\mathrm{MAP}}$ and $\widehat{\boldsymbol{x}}_{\mathrm{MMSE}}$ in terms of MSE (first column) and Jaccard error (second column), while benefiting of significantly lower computational costs. Interestingly, when p increases (large number of change-points), the larger the gain in using the proposed procedure. This is also the case when the sample size N increases: For $N = 10^4$, the MCMC approach takes more than an hour while the method we propose here provides a relevant solution in a few minutes.

We further compared the performance of the proposed strategy with several classical parameter choice based on information criteria such as AIC, SIC, AICC and SICC (see [45] for details about these criteria). Fig. 5 (solid blue) only reports the SICC-based solutions, which performed best among those four criteria. The proposed solution always perform better. Further criteria can be encountered in [46–48]. Finally, we follow

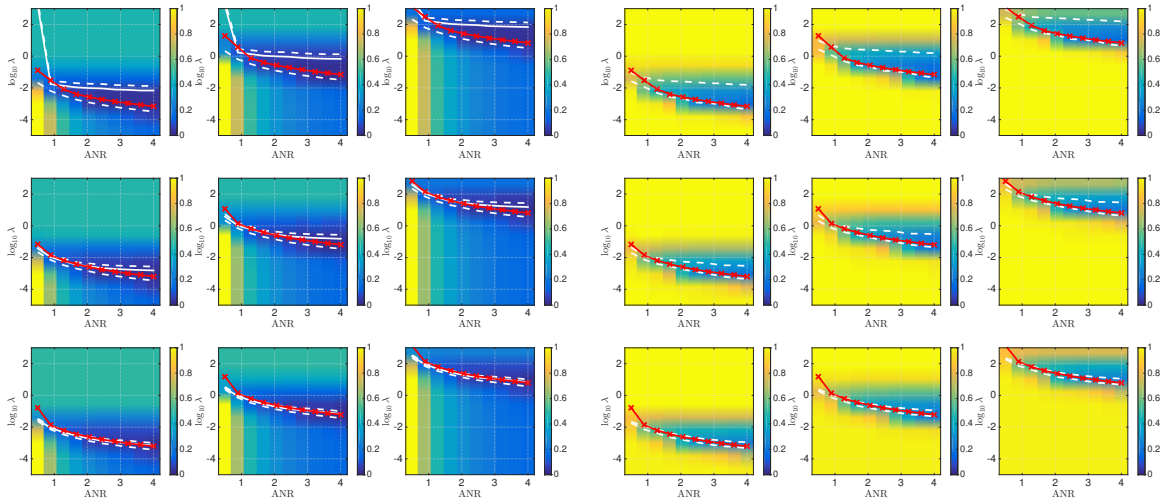

(a) relative MSE (b) Jaccard error

Figure 3: **Estimation performance: RMSE and Jaccard error as functions of λ and** ANR**.** The background displays the relative MSE (left) or Jaccard error (right) w.r.t. λ and ANR. We superimpose in red the estimate $\widehat{\lambda}$ (average over 50 realizations), as a function of the ANR, which is shown to satisfactorily remain within the range of oracles λ, delimited by dashed white lines and to closely follow oracle Monte Carlo average indicated by solid white lines (left: relative MSE, right: Jaccard error). From top to bottom: $p = 0.005$, 0.010 and 0.015. From left to right: $\overline{x}_{\max} - \overline{x}_{\min} = 0.1$, 1, and 10.

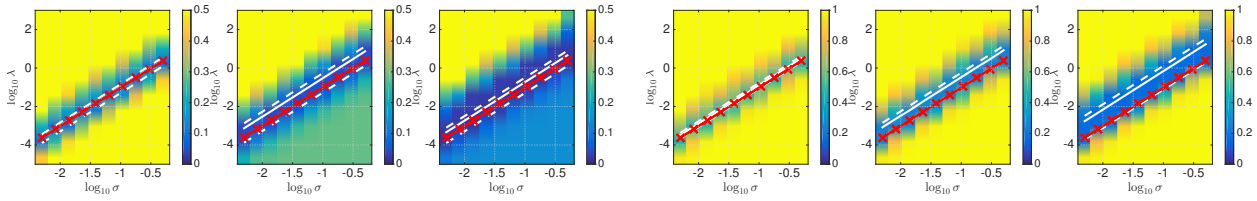

(a) relative MSE (b) Jaccard error

Figure 4: **Estimation performance: RMSE and Jaccard error as functions of λ and σ.** The background displays the relative MSE (left) or Jaccard error (right) w.r.t. λ and σ. We superimpose in red the estimate $\widehat{\lambda}$ (average over 50 realizations), as a function of $\log_{10}\sigma$, which is shown to satisfactorily remain within the range of oracles λ, delimited by dashed white lines and to closely follow oracle Monte Carlo average indicated by solid white lines (left: relative MSE, right: Jaccard error). For each configuration $p = 0.01$ and from left to right: ANR $= 1$, 2 and 4. This illustrates that $\widehat{\lambda}$ leads to solutions with same performance as oracle λ and highlights that $\widehat{\lambda}$ varies linearly with σ^2, as expected.

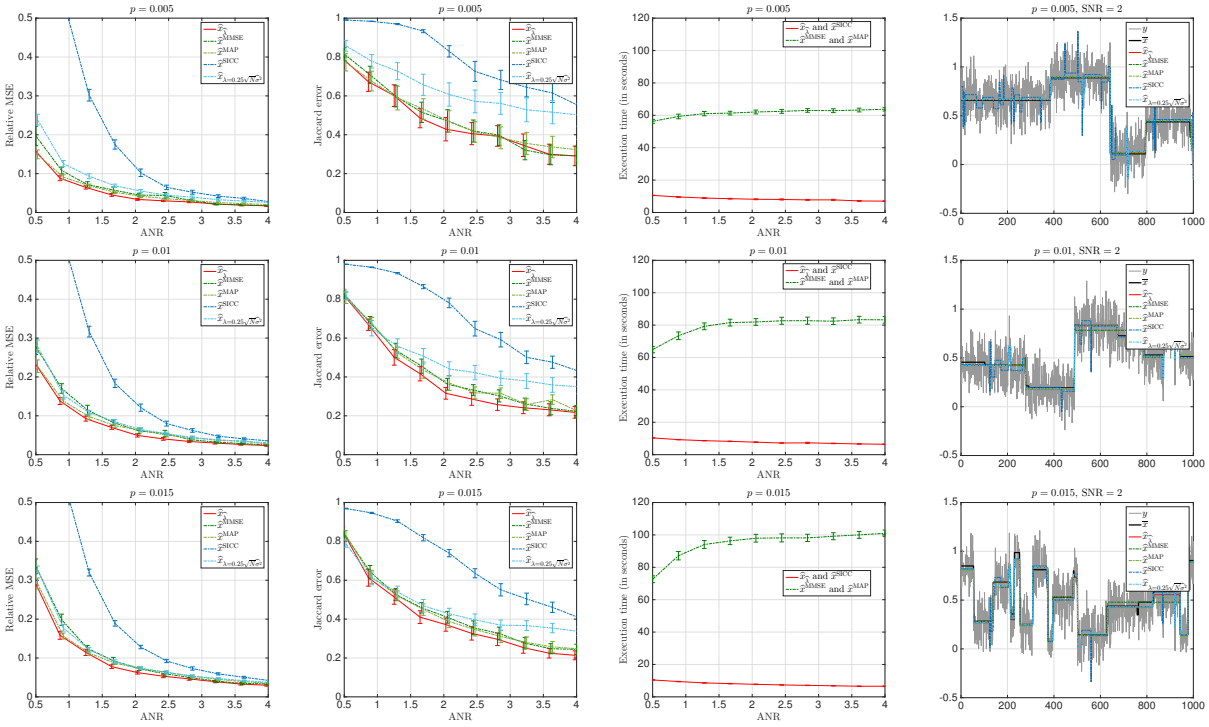

Figure 5: **Comparison with state-of-the-art methods.** For each configuration ANR $= 2$, $\overline{x}_{\max} - \overline{x}_{\min} = 1$ and from top to bottom: $p = 0.005$, 0.010 and 0.015. From left to right: relative MSE, Jaccard error, execution time and example of solutions. The proposed estimator (red) yields estimation performance comparable to Bayesian estimators (green) while benefiting from significantly lower computational costs. Moreover, it improves significantly the performance compared to SICC estimator (blue) for a similar computation cost. The results obtained with $\lambda = 0.25\sqrt{N}\widehat{\sigma}^2$ is displayed in light blue.

the regularization parameter choice provided in [25] consisting in the heuristic rule $\tau = 0.25\sqrt{N}\sigma$ derived in [8] for the ℓ_1-penalized formulation. Invoking dimensionality arguments, this choice as been adapted to $\lambda = 0.25\sqrt{N}\sigma^2$ for the ℓ_0-penalized formulation addressed in this work. One should note that this parameter selection method relies on the perfect knowledge of the noise variance σ^2, which is not the case in the considered study framework. To provide fair comparisons, the performance of this heuristic rule has been evaluated using an estimate of this variance, derived from the classical median estimator [49] which is particularly suitable for piecewise constant signals. Again, the proposed method always lead to better solutions.

5.5 Selection of hyperparameter σ_0^2

We finally investigate the impact of the choice of the hyperparameter σ_0^2 on the achieved solution, according to the discussion in Remark 1. Fig. 6 displays $\widehat{\lambda}$ (red circle) as a function of σ_0^2, for different values of $\overline{x}_{\max} - \overline{x}_{\min}$. It shows that using $\log 2\pi\sigma_0^2 \in [0, 5]$ systematically leads to satisfactory estimates that minimize the relative MSE (left) or Jaccard error (right). This clearly indicates that σ_0^2 does not depend on data dynamics ($\overline{x}_{\max} - \overline{x}_{\min}$), which is what is expected from a hyperparameter. Finally, to explore the potential dependencies on p or ANR, we set $\overline{x}_{\max} - \overline{x}_{\min} = 1$. Fig. 7 shows that selecting any value of σ_0 such that $\log 2\pi\sigma_0^2 \in [0, 5]$ leads to satisfactory estimates minimizing the relative MSE (left) or Jaccard error (right), irrespectively of the actual values of p or of the ANR.

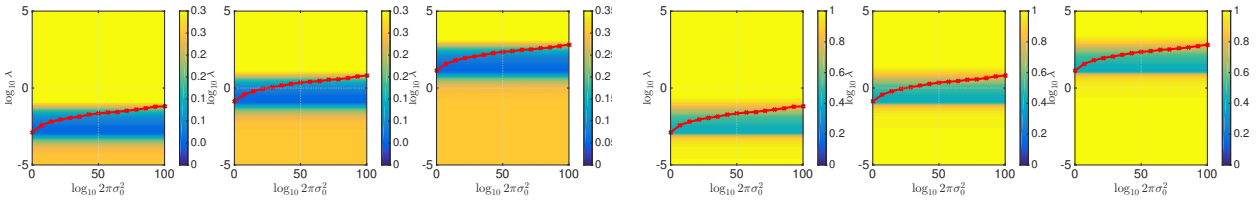

(a) relative MSE (b) Jaccard error

Figure 6: **Estimation performance: RMSE and Jaccard error as functions of λ and σ_0^2.** The background displays the relative MSE (left) or Jaccard error (right) w.r.t. λ and σ_0^2. We superimpose in red the estimate $\widehat{\lambda}$, which a priori explicitly depends on the choice of the hyperprior σ_0^2, is averaged over 50 realizations and displayed in red as a function of $2\pi\sigma_0^2$. From left to right: $\overline{x}_{\max} - \overline{x}_{\min} = 0.1$, 1 and 10. For each configuration $p = 0.01$, ANR $= 1$. Choosing $\log 2\pi\sigma_0^2 \in [0, 5]$ systematically leads to satisfactory estimates minimizing the relative MSE (left) or Jaccard error (right).

6 Conclusions and perspectives

This contribution studied a change-point detection strategy based on the ℓ_2-Potts model, whose performance depend crucially on the selection of a regularization parameter. Using an equivalence between a variational formulation and a hierarchical Bayesian formulation of the change-point detection problem, the present contribution proposed and assessed an efficient automated selection of this regularization parameter. It shows that estimation performance of the proposed procedure (evaluated in terms of global MSE and Jaccard error) match satisfactorily those achieved with oracle solutions. Moreover, when compared to fully Bayesian strategies, the proposed procedure achieved equivalent performance at significantly lower computational costs. One of the advantages of the proposed approach is that it can be easily adapted to different additive noise degradations. For instance, for the ℓ_1-Potts setup, that is the Laplacian noise assumption, the likelihood (11) should be replaced by a Laplacian distribution with scale parameter σ and the step (25) should be replaced by $\widehat{\sigma}_\lambda = \|\boldsymbol{y} - \widehat{\boldsymbol{x}}_\lambda\|/N$ to be consistent. One difficulty that can be encountered for this kind of degradation is that we do not know the conjugate prior for the Laplace distribution. We have recently studied a related issue in a conference paper [50]. Future work could also aim to extend the present framework to Poisson noise.

A Bayesian estimators

The maximum a posteriori (MAP) or minimum mean squared error (MMSE) estimators associated with the joint posterior $f(\boldsymbol{\Theta}|\mathbf{y})$ in (16) can be approximated by using MCMC procedures that essentially rely on a partially collapsed Gibbs sampler [51] similar to the algorithm derived in [5]. It consists in iteratively drawing samples (denoted $\cdot^{[t]}$) according to conditional posterior distributions that are associated with the joint posterior (16). The resulting procedure, detailed in Algo. 2, provides a set of samples $\boldsymbol{\vartheta} = \left\{\boldsymbol{r}^{[t]}, \boldsymbol{\mu}^{[t]}, \sigma^{2[t]}, p^{[t]}\right\}_{t=1}^{T_{\mathrm{MC}}}$ that are asymptotically distributed according to (16). These samples can be used to approximate the MAP and MMSE estimators of the parameters of interest [52]. The corresponding solutions are referred to as $\widehat{\boldsymbol{x}}_{\mathrm{MAP}}$ and $\widehat{\boldsymbol{x}}_{\mathrm{MMSE}}$ in Section 5.4.

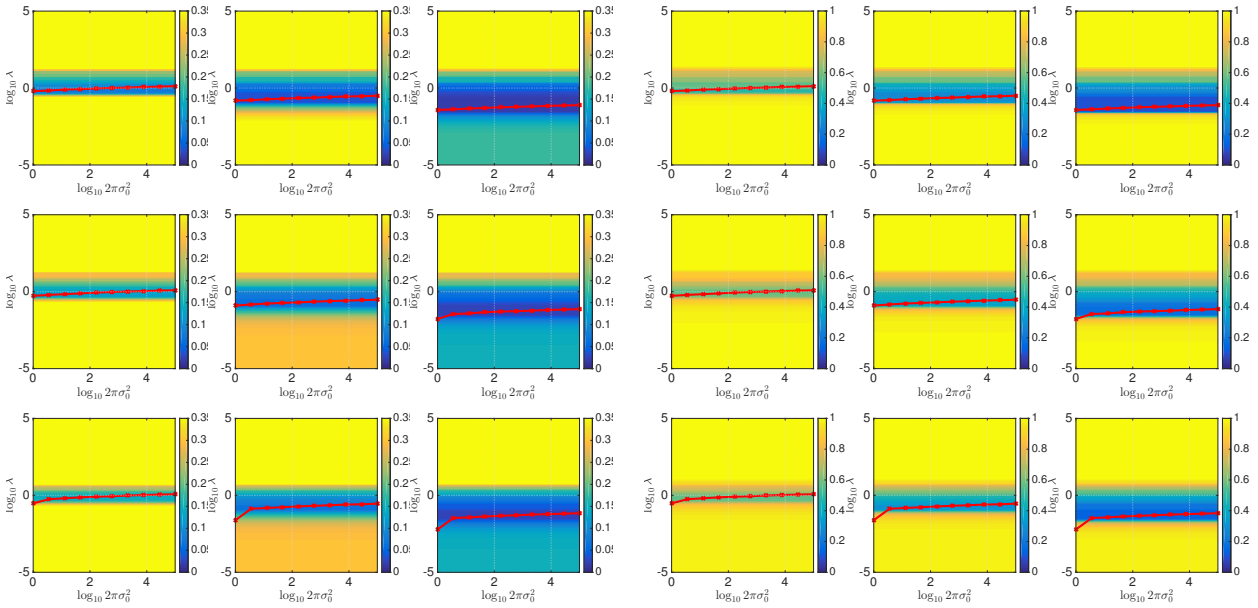

(a) relative MSE (b) Jaccard error

Figure 7: **Estimation performance: RMSE and Jaccard error as functions of λ and σ_0^2.** The background displays the relative MSE (left) or Jaccard error (right) w.r.t. λ and σ_0^2. We superimpose in red the estimate $\widehat{\lambda}$, which a priori explicitly depends on the choice of the hyperprior σ_0^2, is averaged over 50 realizations and displayed in red as a function of $2\pi\sigma_0^2$. Choosing $\log 2\pi\sigma_0^2 \in [0, 5]$ leads to satisfactory estimation performance independently of p and the ANR. For each configuration $\overline{x}_{\max} - \overline{x}_{\min} = 1$. From top to bottom: $p = 0.005$, 0.010 and 0.015. From left to right: ANR = 1, 2 and 4.

Algorithm 2 Piecewise constant Bayesian denoising

Input: Observed signal $\boldsymbol{y} \in \mathbb{R}^N$.
Hyperparameters $\boldsymbol{\Phi} = \{\alpha_0, \alpha_1, \mu_0, \sigma_0^2\}$.

Iterations:

for $t = 1, \ldots, T_{\mathrm{MC}}$ **do**
  for $i = 1, \ldots, N-1$ **do**
    Draw $r_i^{[t]} \sim f\left(r_i | \boldsymbol{y}, \boldsymbol{r}_{\setminus i}, p, \sigma^2, \mu_0, \sigma_0^2\right)$
  end for
  for $k = 1, \ldots, \sum_{i=1}^N r_i^{[t]}$ **do**
    Draw $\mu_k^{[t]} \sim f\left(\mu_k | \boldsymbol{y}, \boldsymbol{r}, \sigma^2, \mu_0, \sigma_0^2\right)$
  end for
  Draw $\sigma^{2[t]} \sim f\left(\sigma^2 | \boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\mu}\right)$
  Draw $p^{[t]} \sim f\left(p | \boldsymbol{r}, \alpha_0, \alpha_1\right)$
end for

Output: $\boldsymbol{\vartheta} = \left\{\boldsymbol{r}^{[t]}, \boldsymbol{\mu}^{[t]}, \sigma^{2[t]}, p^{[t]}\right\}_{t=1}^{T_{\mathrm{MC}}}$, $\widehat{\boldsymbol{x}}_{\mathrm{MAP}}$ and $\widehat{\boldsymbol{x}}_{\mathrm{MMSE}}$.

References

[1] M. Basseville and I.V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.

[2] M.A. Little and N.S. Jones, "Generalized methods and solvers for noise removal from piecewise constant signals. I. Background theory," *Proc. R. Soc. A*, vol. 467, pp. 3088–3114, 2011.

[3] Y. Sowa, A.D. Rowe, M.C. Leake, T. Yakushi, M. Homma, and R.M. Ishijima, A. Berry, "Direct observation of steps in rotation of the bacterial flagellar motor," *Nature*, vol. 437, pp. 916–919, Oct. 2005.

[4] M. Lavielle and E. Lebarbier, "An application of MCMC methods for the multiple change-points problem," *Signal Process.*, vol. 81, no. 1, pp. 39–53, Jan. 2001.

[5] N. Dobigeon, J.-Y. Tourneret, and J. D. Scargle, "Joint segmentation of multivariate astronomical time series: Bayesian sampling with a hierarchical model," *IEEE Trans. Signal Process.*, vol. 55, no. 2, pp. 414–423, Feb. 2007.

[6] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys.D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.

[7] S.-J. Kim, K. Koh, S. Boyd, and D. Gorinevsky, "ℓ_1-trend filtering," *SIAM Review*, vol. 51, no. 2, pp. 339–360, May 2009.

[8] L. Dümbgen and Kovac A., "Extensions of smoothing via taut strings," *Electronic Journal of Statistics*, vol. 3, pp. 41–75, 2009.

[9] M. A. Little and N. S. Jones, "Sparse bayesian step-filtering for high-throughput analysis of molecular machine dynamics," in *Proc. IEEE Int. Conf. Acoust., Speech and Signal Process. (ICASSP)*, Dallas, TX, March 2010, pp. 4162–4165.

[10] I. W. Selesnick and I. Bayram, "Total variation filtering," Tech. Rep., 2010.

[11] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the bayesian restoration of images," *IEEE Trans. Pattern Anal. Match. Int.*, vol. 6, no. 6, pp. 721–741, Jun. 1984.

[12] Y.-C. Yao, "Estimation of a noisy discrete-time step function: Bayes and empirical Bayes approaches," *Ann. Statist.*, vol. 12, no. 4, pp. 1434–1447, 1984.

[13] D. Mumford and J. Shah, "Optimal approximations by piecewise smooth functions and associated variational problems.," *Comm. Pure Applied Math.*, vol. 42, pp. 577–685, 1989.

[14] G. Winkler and V. Liebscher, "Smoothers for discontinuous signals," *Journal of Nonparametric Statistics*, vol. 14, no. 1-2, pp. 203–222, 2002.

[15] C.R. Vogel and M.E. Oman, "Iterative methods for total variation denoising," *J. Sci. Comput.*, vol. 17, pp. 227–238, 1996.

[16] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, no. 1-2, pp. 89–97, Jan. 2004.

[17] A. Barbero and S. Sra, "Fast Newton-type methods for total variation regularization," in *International Conference on Machine Learning*, Lise Getoor and Tobias Scheffer, Eds., New York, NY, USA, Jun. 2011, ICML '11, pp. 313–320, ACM.

[18] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, New York, 2011.

[19] L. Condat, "A direct algorithm for 1D total variation denoising," *IEEE Signal Process. Lett.*, vol. 20, no. 11, pp. 1054–1057, 2013.

[20] P.L. Davies and A. Kovac, "Local extremes, runs, strings and multiresolution," *Ann. Stat.*, vol. 29, no. 1, pp. 1–65, 2001.

[21] C.M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Stat.*, vol. 9, no. 6, pp. 1135–1151, 1981.

[22] C. Deledalle, S. Vaiter, G. Peyré, J.M. Fadili, and C. Dossal, "Unbiased risk estimation for sparse analysis regularization," in *Proc. Int. Conf. Image Process. (ICIP)*, Orlando, FL, USA, Sept. 30-Oct. 3 2012.

[23] A. Benazza-Benyahia and J.-C Pesquet, "Building robust wavelet estimators for multicomponent images using Stein's principles," *IEEE Trans. Image Process.*, vol. 14, no. 11, pp. 1814–1830, Nov. 2005.

[24] J.-C. Pesquet, A. Benazza-Benyahia, and C. Chaux, "A SURE approach for digital signal/image deconvolution problems," *IEEE Trans. Signal Process.*, vol. 57, pp. 4616–4632, 2009.

[25] I. W. Selesnick, A. Parekh, and I. Bayram, "Convex 1-D total variation denoising with non-convex regularization,"," *IEEE Signal Process. Lett.*, vol. 22, no. 2, pp. 141–144, Feb. 2015.

[26] J. P. Oliveira, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Adaptive total variation image deblurring: a majorization–minimization approach," *Signal Process.*, vol. 89, no. 9, pp. 1683–1693, 2009.

[27] L. Chaari, J.-C. Pesquet, J.-Y. Tourneret, and P. Ciuciu, "Parameter estimation for hybrid wavelet-total variation regularization," in *Proc. IEEE Workshop Stat. Sign. Proc.*, Nice, France, June, 28-30 2011.

[28] M. Pereyra, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Maximum-a-posteriori estimation with unknown regularisation parameters," in *Proc. Eur. Sig. Proc. Conference*, Nice, France, Aug 2015, pp. 230–234.

[29] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, "Estimating the granularity coefficient of a Potts-Markov random field within an MCMC algorithm," *IEEE Trans. Image Process.*, vol. 22, no. 6, pp. 2385–2397, June 2013.

[30] G. Winkler, A. Kempe, V. Liebscher, and O. Wittich, *Innovations in Classification, Data Science, and Information Systems: Proceedings of the 27th Annual Conference of the Gesellschaft für Klassifikation e.V., Brandenburg University of Technology, Cottbus, March 12–14, 2003*, chapter Parsimonious Segmentation of Time Series by Potts Models, pp. 295–302, Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

[31] F. Friedrich, A. Kempe, V. Liebscher, and G. Winkler, "Complexity penalized M-estimation: fast computation," *J. Comput. Graph. Statist.*, vol. 17, no. 1, pp. 201–224, 2008.

[32] M. Storath, A. Weinmann, and L. Demaret, "Jump-sparse and sparse recovery using Potts functionals," *IEEE Trans. Signal Process.*, vol. 62, no. 14, pp. 3654–3666, July 2014.

[33] M. Storath, A. Weinmann, and M. Unser, "Exact algorithms for l^1-TV regularization of real-valued or circle-valued signals," *J. Sci. Comput.*, vol. 38, no. 1, pp. 614–630, 2016.

[34] M. Lavielle, "Optimal segmentation of random processes," *IEEE Trans. Signal Process.*, vol. 46, no. 5, pp. 1365–1373, 1998.

[35] N. Dobigeon, J.-Y. Tourneret, and M. Davy, "Joint segmentation of piecewise constant autoregressive processes by using a hierarchical model and a Bayesian sampling approach," *IEEE Trans. Signal Process.*, vol. 55, no. 4, pp. 1251–1263, Apr. 2007.

[36] N. Dobigeon and J.-Y. Tourneret, "Joint segmentation of wind speed and direction using a hierarchical model," *Comput. Stat. Data Anal.*, vol. 51, no. 12, pp. 5603–5621, Aug. 2007.

[37] E. Punskaya, C. Andrieu, A. Doucet, and W. Fitzgerald, "Bayesian curve fitting using MCMC with applications to signal segmentation," *IEEE Trans. Signal Process.*, vol. 50, no. 3, pp. 747–758, Mar. 2002.

[38] G. Kail, J.-Y. Tourneret, F. Hlawatsch, and N. Dobigeon, "Blind deconvolution of sparse pulse sequences under a minimum distance constraint: a partially collapsed Gibbs sampler method," *IEEE Trans. Signal Processing*, vol. 60, no. 6, pp. 2727–2743, June 2012.

[39] N. Dobigeon, J.-Y. Tourneret, and J. D. Scargle, "Joint segmentation of multivariate poissonian time series applications to burst and transient source experiments," in *Proc. European Signal Processing Conf. (EUSIPCO)*, Florence, Italy, Sept. 2006.

[40] Harold Jeffreys, "An invariant form for the prior probability in estimation problems," *Proc. of the Royal Society of London. Series A*, vol. 186, no. 1007, pp. 453–461, 1946.

[41] Christian P. Robert, *The Bayesian Choice: from Decision-Theoretic Motivations to Computational Implementation*, Springer Texts in Statistics. Springer-Verlag, New York, 2 edition, 2007.

[42] M. Storath and A. Weinmann, "Fast partitioning of vector-valued images," *SIAM J. Imaging Sci.*, vol. 7, no. 3, pp. 1826–1852, 2014.

[43] P. Jaccard, "Distribution de la flore alpine dans le bassin des Dranses et dans quelques régions voisines," *Bulletin de la Société Vaudoise des Sciences Naturelles*, , no. 37, pp. 241–272, 1901.

[44] R. Hamon, P. Borgnat, P. Flandrin, and C. Robardet, "Relabelling vertices according to the network structure by minimizing the cyclic bandwidth sum," *Journal of Complex Networks*, vol. 4, no. 4, pp. 534–560, 2016, Accepted to *Journal of Complex Networks*, *arXiv preprint arXiv:1410.6108*.

[45] G. Winkler, O. Wittich, V. Liebscher, and A. Kempe, "Don't shed tears over breaks," *Jahresber. Deutsch. Math.-Verein*, vol. 107, pp. 57–87, 2005.

[46] N.R. Zhang and D.O. Siegmund, "A modified Bayes information criterion with applications to the analysis of comparative genomic hybridization data," *Biometrics*, vol. 63, pp. 22–32, Mar. 2007.

[47] K. Frick, A. Munk, and H. Sieling, "Multiscale change point inference.," *Journal of the Royal Statistical Society:Series B*, vol. 76, no. 3, pp. 495–580, 2014.

[48] Z. Xia and P. Qiu, "Jump information criterion for statistical inference in estimating discontinuous curves," *Biometrika*, vol. 102, no. 2, pp. 397–408, 2015.

[49] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. American Statist. Ass.*, vol. 90, no. 432, pp. 1200–1224, Dec. 1995.

[50] J. Frecon, N. Pustelnik, N. Dobigeon, H. Wendt, and P. Abry, "Bayesian-driven criterion to automatically select the regularization parameter in the l1-Potts model," New Orleans, USA, Mar. 5-9 2017.

[51] D.A. van Dyk and T. Park, "Partially collapsed Gibbs samplers: Theory and methods," *J. American Statist. Ass.*, vol. 103, no. 482, pp. 790–796, June 2008.

[52] J.M. Marin and C.P. Robert, *Bayesian Core: A Practical Approach to Computational Bayesian Statistics*, Springer, New York, NY, USA, 2007.